DUMPSARENA

NetApp Accredited Hardware Support Engineer

Netapp NS0-093

Version Demo

Total Demo Questions: 8

Total Premium Questions: 82

Buy Premium PDF

https://dumpsarena.co

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.co

Topic Break Down

Topic	No. of Questions
Topic 1, NetApp Storage System Architecture	17
Topic 2, Hardware Installation	2
Topic 3, Hardware Maintenance	29
Topic 4, Hardware Troubleshooting	34
Total	82

QUESTION NO: 1

A node has rebooted after a panic, and the node is now reachable through the cluster CLI.

Which two actions help preserve diagnostic information for NetApp Support analysis? (Choose two.)

A. Delete the core dump to release root-volume space.

B. Run `system node coredump show`.

C. Reboot the node again before reviewing its diagnostic files.

D. Run `system node autosupport invoke-core-upload`.

ANSWER: B D

Explanation:

`system node coredump show` identifies the core dump generated by the panic and confirms whether the diagnostic artifact is available. The core dump contains kernel state captured at the time of the failure. The `system node autosupport invoke-core-upload` command initiates upload of the retained core file through AutoSupport when the required transport is configured.

Deleting a core dump removes evidence needed for investigation. Rebooting the node again does not collect the existing panic data and can complicate analysis of the original event.

QUESTION NO: 2

What are two valid options for uploading a core file from a node that is running ONTAP 9.12.1 software to NetApp for analysis? (Choose two.)

A. system node autosupport invoke -diagnostic

B. Download via CIFS and upload to upload.netapp.com.

C. Download via Service Processor Infrastructure (SPI) and upload to upload.netapp.com.

D. system node autosupport invoke-core-upload

ANSWER: B D

Explanation:

Download via CIFS and upload to upload.netapp.com. is a valid manual workflow for getting a node core file to NetApp Support. An administrator can retrieve the core file from the node through a CIFS-accessible location, then submit the downloaded file through NetApp's secure upload service. This is useful when direct AutoSupport delivery is unavailable or when Support requests a specific core file.

system node autosupport invoke-core-upload is the supported ONTAP command for initiating the upload of a core file through the configured AutoSupport transport. The command submits the selected core dump to NetApp, avoiding the need to first copy the potentially large diagnostic file to an administrative workstation. Before using this method, the cluster must have a working AutoSupport delivery configuration and the administrator must identify the applicable node and core file as required by the command syntax.

Core files contain diagnostic state captured during a node failure or panic and are used by NetApp Support to investigate the event. NetApp documents core-dump administration and the available upload workflow in the ONTAP documentation: Manage core dumps and system node autosupport invoke-core-upload command reference.

QUESTION NO: 3

What is the recommended value for disk and CPU use when you plan an upgrade?

A. less than 50%

B. less than 90%

C. less than 85%

D. less than 70%

ANSWER: A

Explanation:

The recommended planning threshold is **less than 50%** for both disk and CPU use. An ONTAP software upgrade performs additional control-plane work, including validation, package staging, takeover/giveback activity in an HA pair, and service restarts. Maintaining at least half of the available CPU and disk-performance capacity as headroom helps the cluster absorb this temporary workload while continuing to serve client I/O. This threshold is therefore a pre-upgrade readiness check, not merely a general performance target. Before beginning the upgrade, the support engineer should review node CPU activity and aggregate/disk workload during a representative busy period, rather than relying only on an idle-time reading. If either measure is near or above the threshold, the upgrade should be scheduled for a quieter interval or the workload should be reduced first. NetApp's ONTAP upgrade preparation guidance emphasizes verifying cluster readiness and system health before initiating the upgrade; preserving resource headroom is part of performing a safe nondisruptive upgrade. See the NetApp ONTAP upgrade preparation documentation and the ONTAP system administration documentation for upgrade-readiness and monitoring guidance.

QUESTION NO: 4

A controller has been repaired after a boot-device replacement and is ready to return to service. Its partner remains in takeover.

Which two checks should be completed before initiating giveback? (Choose two.)

A. Verify that the repaired node has completed booting and is ready for giveback.

B. Verify the HA relationship and takeover status.

C. Power off the HA partner before starting giveback.

D. Boot the repaired node into maintenance mode for normal client I/O.

ANSWER: A B

Explanation:

The repaired node must complete its boot and reach a state in which it is ready to receive its aggregates. The engineer should also verify HA status to confirm that the partner relationship is healthy before returning aggregate ownership. A giveback should not be initiated while the node is still in the LOADER environment or while an unresolved hardware fault remains.

QUESTION NO: 5

Which type of core file is generated when a node panics?

A. mgwd core

B. user space core

C. sync core

D. kernel core

ANSWER: D

Explanation:

kernel core is generated when an ONTAP node panics. A panic is a critical kernel-level failure that causes ONTAP to halt or restart the node to protect data integrity. During this event, ONTAP captures a kernel core dump: a diagnostic image containing the kernel's state at the time of failure, including relevant memory and execution information. NetApp Support uses this artifact to identify the fault path and determine the underlying cause of the panic. After reboot, ONTAP retains core-dump information so that it can be reviewed and, when required, sent to NetApp. Administrators can view available core dumps with the `system node coredump show` command and can initiate an AutoSupport core upload with `system node autosupport invoke-core-upload`. Timely retention and upload of the kernel core is important because it provides the detailed crash evidence required for effective root-cause analysis. NetApp documents panic recovery and core-dump collection in its ONTAP system-administration guidance and CLI command references: ONTAP node boot and recovery documentation and system node coredump show command reference.

QUESTION NO: 6

In a SAS stack of shelves, what is the topology of the connection between expander and disk?

A. arbritrated loop

B. point-to-point

C. loop

D. ring

ANSWER: B

Explanation:

point-to-point is correct because Serial Attached SCSI uses serial links between individual SAS phys. Within a NetApp SAS disk shelf, each drive port connects through its own dedicated SAS physical link to a PHY on the shelf's I/O module (the SAS expander). The expander then provides the switching function that permits initiators and other expanders to reach the individual disk devices. Therefore, the direct expander-to-disk relationship is point-to-point, even though the overall shelf stack can contain multiple devices and redundant paths.

This architecture is fundamental to SAS operation: a SAS domain is formed from end devices and expanders connected by point-to-point links. Disk shelves use expanders to scale the number of attached drives while retaining independent physical connections for the drives. NetApp's platform installation documentation illustrates SAS shelf cabling and the redundant, multipath shelf connectivity used by supported systems; the underlying disk-to-expander links inside each shelf remain SAS point-to-point links. See the NetApp AFF C800 SAS shelf cabling documentation and the T10 SCSI standards organization for SAS standards information.

QUESTION NO: 7

A temporary shelf connectivity failure caused several healthy disks to be marked as failed. The shelf I/O module and cabling fault has been corrected.

Which two conditions should be confirmed before you unfail the disks? (Choose two.)

A. The shelf connectivity fault has been corrected and the disks are consistently visible.

B. The node boot media has been replaced with a new device.

C. The disks show no persistent media or hardware errors.

D. The shelf has been power cycled after the I/O module replacement.

ANSWER: A C

Explanation:

The original connectivity fault must be corrected and the disks must be visible consistently before an unfail operation is considered. Unfailing a disk only removes the failed marking; it does not repair an unresolved SAS path, IOM, or media problem. The engineer must also confirm that the affected disks do not show persistent disk or media errors that would make returning them to service unsafe.

Updating ONTAP or replacing boot media does not establish that the affected disks and their shelf paths are healthy. A shelf power cycle can introduce an unnecessary disruption and is not a prerequisite for unfailing disks.

QUESTION NO: 8

After a SAS HBA replacement, a controller can no longer see any disks in one shelf stack. Other shelf stacks connected to the controller remain visible.

What should be investigated first?

A. Disk ownership for every disk in the affected shelf.

B. The boot device installed in the controller.

C. The SAS cable routing and HBA port connections for the affected stack.

D. The WAFL consistency state of the root aggregate.

ANSWER: C

Explanation:

The SAS cabling and port placement for the affected stack should be checked first against the applicable cabling worksheet or Hardware Universe configuration. The fact that other shelf stacks remain visible narrows the issue to the newly serviced HBA path or the physical connection to the affected stack. A cable in an incorrect port, an incompletely seated cable, or an unsupported port arrangement can make an entire stack unavailable.

Disk ownership does not normally cause every disk in one connected stack to disappear. Replacing the boot device or running WAFLiron would not restore physical SAS discovery.